

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 November 2000 (30.11.2000)

PCT

(10) International Publication Number
WO 00/72532 A1

(51) International Patent Classification⁷: H04L 12/56, 29/06

(21) International Application Number: PCT/US00/14381

(22) International Filing Date: 24 May 2000 (24.05.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/135,605 24 May 1999 (24.05.1999) US

(71) Applicant: RUTGERS, THE STATE UNIVERSITY OF NEW JERSEY [US/US]; Old Queens Building, Somerset and George Streets, New Brunswick, NJ 08903 (US).

(71) Applicants and

(72) Inventors: BADRINATH, B.R. [US/US]; 4 Depalma Court, Somerset, NJ 08873 (US). SUDAME, Predeep [US/US]; 18J Rele Lane, Somerset, NJ 08873 (US).

(74) Agent: VIKSNINS, Ann, S.; Schwegman, Lundberg, Woessner & Kluth, P.O. Box 2938, Minneapolis, MN 55402 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

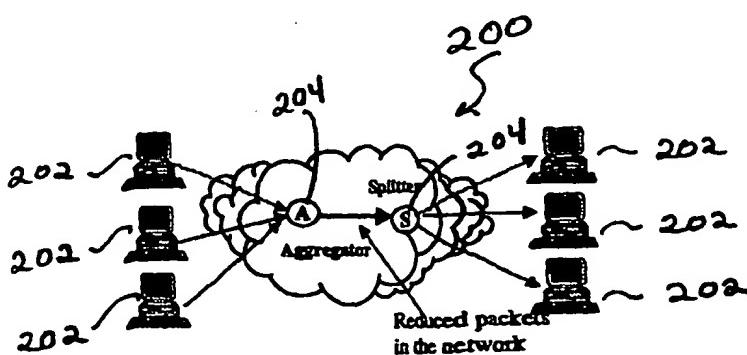
(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

- With international search report.
- Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYSTEM AND METHOD FOR NETWORK PACKET REDUCTION



(57) Abstract: A system of and method for reducing network packet traffic is described. Embodiments of the present invention combine small packets with the same destination into a larger packet, thereby reducing the number of packets flowing in the network. Small packets dominate Internet traffic patterns today. The overhead due to queuing, processing, routing, differential treatment, and delivery is the same irrespective of the size of the packet. Hence, small packets impose a relatively higher overhead on the network. The domination of small packet traffic is expected to get only worse with the ever

WO 00/72532 A1

increasing web applications and the introduction of new applications such as Voiceover IP.

BEST AVAILABLE COPY

System and Method for Network Packet Reduction

5

Field of the Invention

The present invention relates generally to computer networks and more specifically to reducing packets on a network.

Copyright Notice/Permission

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever. The following notice applies to the software and data as described below and in the drawings hereto: Copyright © 2000 Rutgers, The State University of New Jersey, All Rights Reserved.

Background of the Invention

Data communication in a computer network involves the exchange of data between two or more entities interconnected by communication links and subnetworks. These entities are typically software programs executing on hardware computer platforms, such as end stations and intermediate stations. Examples of an intermediate station may be a router or switch which interconnects the communication links and subnetworks to enable transmission of data between the end stations. A local area network (LAN) is an example of a subnetwork that provides relatively short distance communication among the interconnected stations; in contrast, a wide area network (WAN) enables long distance communication over links provided by public or private telecommunications facilities. Communication software executing on the end stations correlates and manages data communication with other end stations. The stations typically communicate by exchanging discrete

packets or frames of data according to predefined protocols.

Collectively, these hardware and software components comprise a communications network.

Small packets (less than the default link MTU of 512 bytes) dominate Internet traffic patterns today. Wide-area Internet traffic measurements show that as many as sixty percent (60%) of the packets transmitted over the Internet backbone are less than 44 bytes. Web servers, for example, are a leading generator of small packets. All of the TCP ACKs generated by the clients have the same destination (the web server). These ACKs are just 40 bytes in size. Another example of small packet traffic is ICP (Internet Cache Protocol) queries generated among web proxies. ICP queries are about 60 bytes in size. This domination of small packets in Internet traffic patterns is expected to get only worse with the ever increasing web traffic and with introduction of new applications such as VoIP (Voiceover-IP) applications generate small audio samples, about 40 to 60 byte packets every 40 ms). The network overhead cost due to queuing, processing, routing, differential treatment and delivery is the same irrespective of the size of the packet. Hence, small packets impose a relatively higher overhead on the network.

Therefore, in order to improve overall network performance, there is a need to reduce the number of packets in the network.

Summary of the Invention

A system of and method for reducing network packet traffic is described. Embodiments of the present invention combine small packets with the same destination into a larger packet, thereby reducing the number of packets flowing in the network. Small packets dominate Internet traffic patterns today. The overhead due to queuing, processing, routing, differential treatment, and delivery is the same irrespective of the size of the packet. Hence, small packets impose a relatively higher overhead on the network. The domination of small packet traffic is expected to get only worse with the ever increasing web applications and

the introduction of new applications such as Voiceover IP. To mitigate the impact of small packets, a novel approach of combining multiple small packets going towards the same destination is provided.

One aspect of the invention is a computerized method of reducing 5 packets on a network. The method comprises combining multiple packets for a destination into an aggregated packet, and separating the aggregated packet into the multiple packets at the destination.

Still other and further embodiments, aspects and advantages of the 10 invention will become apparent by reference to the drawings and by reading the following detailed description.

Brief Description of the Drawings

Figure 1 is a block diagram of a traditional communication network.

Figure 2 is a block diagram of a communication network according to 15 an example embodiment of the invention.

Figure 3A is a block diagram of a traditional router.

Figure 3B is a block diagram of an aggregator, such as the aggregator shown in Figure 2, according to an example embodiment of the present 20 invention.

Figure 4 is a block diagram of a splitter, such as the splitter shown in 25 Figure 2, according to an example embodiment of the present invention.

Figure 5 is a block diagram showing how metadata is stored when two small packets are combined at a router according to an example 30 embodiment of the invention.

Description of the Embodiments

In the following detailed description of the embodiments, reference is 25 made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention. It is to be understood that other embodiments may be utilized and that 30 structural, logical and electrical changes may be made without departing

from the scope of the present inventions. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present inventions is defined only by the appended claims.

A system of and method for reducing network packet traffic is described below. Embodiments of the present invention combine small packets with the same destination into a larger packet, thereby reducing the number of packets flowing in the network.

Figure 1 is a block diagram of a traditional communication network 100. As shown in Figure 1, the traditional communication network 100 comprises a plurality of end stations 102 and a plurality of intermediate stations 104 interconnected by communications links. In the example shown in Figure 1, the intermediate links are routers 104.

Figure 2 is a block diagram of a communication network 200 according to an example embodiment of the present invention. As shown in Figure 2, the communication network 200 comprises a plurality of end stations 202 and a plurality of intermediate stations 204. The intermediate stations 204 include an aggregator and a splitter which are further described below. The aggregator and the splitter function to reduce the packets on the communications link between them. The aggregator combines small packets into a larger packet so that the routers see a smaller number of packets. The splitter separates these packets so that the subsequent stations in the network are unaware of the packet aggregation. Aggregators, as well as splitters, can be placed at intermediate stations or at end stations in the network. The placement of aggregators and splitters may be varied depending on traffic patterns, level of congestion and software application needs.

In one embodiment, the aggregators as well as the splitters can be the end stations in the network. In an alternate embodiment, the aggregators and splitters are placed at ingress and egress routers in the network. The placement of aggregators and splitters may be manually configured or may be dynamically assigned according to a signaling mechanism of the

present invention which is further described below. In another embodiment, the aggregated packets on the communications links between the aggregator and the splitter are treated as a special class of traffic to further enhance network performance. The number of packets in the network are reduced by combining small packets multiple times at various points in the network.

In one embodiment, the placement of the aggregator in the network is based on the application and is dynamically determined by the receiver of the packets. A signaling mechanism of the present invention allows the receiver of the packets to place aggregators at points in the network where packets can be combined. If the aggregation at that location in the network is ineffective, the aggregators can be dynamically or manually removed and, if desired, repositioned at another point in the network. In other words, if the application ceases to need aggregation, the network can remove the aggregators without affecting the normal network routing. Thus, the signalling mechanism allows the aggregation of packets to be initiated and stopped dynamically. In an alternate embodiment, the signaling mechanism can be initiated and stopped by manual configure according to a user.

Figure 3A is a block diagram of a traditional router 300. As shown in Figure 3A, three packets 302 are received by router 300. The router 300 performs a routing table 304 lookup to determine the output queue 306 for each one of the packets. All three packets 302 leave router 300 and are transmitted across a communication link.

Figure 3B is a block diagram of an aggregator such as the aggregator shown in Figure 2. As used herein, an aggregator refers to any end station or intermediate station capable of combining multiple small packets to the same destination and sending them as one packet. For example, an aggregator may be implemented as a hardware module or a software module in a router. As shown in Figure 3B, three packets 312 are received by router 310. The router 310 performs a routing table 304 lookup to determine the destination of the packets

302. However, unlike the traditional router shown in Figure 3A, the router 310 comprises an aggregation module 315. In this example, the three packets 302 have the same destination. The aggregation module 315 combines the three packets 312 into a single packet 318 referred to herein as an "aggregated packet." 5 The aggregated packet 318 is placed in the output queue 316 for the destination and unlike the traditional router, only the single aggregated packet leaves the router 310 and is transmitted across the communication link.

Once such an aggregated packet reaches the destination, all the original packets can be regenerated. The node that separates aggregated packets is 10 referred to herein as a "splitter." Figure 4 is a block diagram of a splitter 400. As used herein, a splitter refers to any end station or intermediate station capable of separating an aggregated packets into the original multiple packets. For example, a splitter may be implemented as a hardware module or a software module in a router. As shown in Figure 4, an aggregated packet 402 is received 15 by router 400. The router 400 performs a routing table 404 lookup to determine the destination of the packet 402. However, unlike the traditional router shown in Figure 3A, the router 400 comprises a splitter module 406. The splitter module 400 separates the aggregated packet 406 into the original multiple packets 408. In one embodiment, the multiple packets 408 are placed in the 20 output queue 410 for router 400.

Aggregation as well as splitting can be performed either at layer 2 (MPLS) or at layer 3 (network layer) of the protocol stack. Aggregators at layer 2 require a hardware implementation. In an example embodiment, aggregators and splitters are implemented using transformer tunnels. Transformer tunnels are 25 described in more detail in Predeep Sudame and B. R. Bardinath, "Transformer tunnels: A framework for providing route-specific adaptions," in *Proceedings of the USENIX Annual Technical Conference*, June 1998, pp. 191-200, which is herein incorporated by reference. However, embodiments of the present invention are not limited to being implemented with transformer tunnels and 30 alternate embodiments are contemplated. To control the number of routing table entries at an aggregator node, the splitter can be a node other than the final

destination. This way, packets belonging to multiple destinations (for example, a subnet) can be combined using a single tunnel. Aggregators may also be dynamically placed in the network so that the bottleneck links lie between the splitters and the aggregators.

5 A computerized method of reducing packets on a network according an example embodiment of the invention comprises combining multiple packets for a destination into an aggregated packet, transmitting the single packet to the destination, and separating the aggregated packet into the multiple packets at the destination. In one embodiment, an aggregation function combines the multiple 10 packets for a destination into an aggregated packet using a mechanism similar to TCP delayed ACKs. Every small packet is delayed by a small amount of time. The process of combining small packets continues until either m (a parameter to the aggregation function described below) packets have been combined, or a timeout occurs. If no other small packets arrive before the timeout, the delayed 15 packet is sent as is. This aggregation mechanism is different from IP reassembly (performed after IP fragmentation). An aggregation function combines small packets, along with their IP headers, to get a larger packet. The receiver regenerates all the original packets when it receives such an aggregated packet. Thus, this mechanism can be used even for protocols that honor message 20 boundaries (e.g., UDP).

In one embodiment, for the receiver to be able to regenerate the original packets, some information is added (metadata) to the aggregated packet. Figure 5 is a block diagram showing how metadata is stored when two small packets are combined at a router. The generated larger packet has a new protocol number 25 (77) in the IP header. (The original protocol therefore has to be stored as part of the metadata, along with the original destination of the packet, if it is different than the address of the splitter node.) The receiver of this packet handles the packet to the protocol handler for this new protocol. This handler is responsible for removing all the metadata and restoring the packets.

30 As small packets are delayed for some time in the hope that they can be combined with other small packets, the delayed packets see increased latency.

The latency depends on the packet inter-arrival rate, and the timeout set by the applications. In the best case, when m packets are combined, the timeout does not come into picture, because the buffers are flushed immediately when the packets are combined.

5 An example embodiment of an aggregation function as described above is provided below. The aggregation function is provided for illustrative purposes only.

Parameters:

10 timeout, m : max packets to be combined

Initialization code:

delayed_packet = null;

15 aggregate (p) {
 if (p is large) {
 cancel timer; send delayed_packet;
 delayed_packet=null;
 }
 20 send p; return PACKET_SENT;
 }
 if (delayed_packet exists) {
 if (p and delayed_packet can be combined) {
 delayed_packet = combine (delayed_packet, p);
 25 if (m packets have been combined) {
 cancel timer; send delayed_packet;
 delayed_packet = null;
 }
 return PACKET_SENT;
 }
 30 } else {
 cancel timer; send delayed_packet;
 set timer; delayed_packet = p;
 return PACKET_SENT;
 }
 35 } else {
 set timer; delayed_packet = p;
 return PACKET_SENT;
 }
 }
40 Timer expires:
 send delayed_packet; delayed_packet = null;

Given that there is sufficient evidence of small packet traffic in the Internet, there are a number of benefits to be obtained by packet-level aggregation. First, most router implementations (such as FreeBSD, Linux, and many Cisco routers) are packet limited rather than byte limited. They place a limit on the number of packets that the router can queue irrespective of the size of each packet. For reasons of efficiency, routers store packet headers (whose sizes are independent of the packet sizes) in fast memory, and the packet payload in slow memory. The fast memory being expensive decides the queue size. Every packet, irrespective of the size, consumes one slot in the queue.

Combining small packets into a larger packet reduces the number of packets seen at routers, and hence helps absorb bursts (reduce losses) for a given level of congestion.

Second, every packet at a router requires a routing table lookup and incurs the same performance cost irrespective of the size of the packet. Combining small packets at a router reduces the number of packets seen by all the subsequent routers. Moreover, to implement a specific per-hop behavior (as in differential services), routers have to give preferential treatment to every packet of a particular class. By combining small packets belonging to the same class, these overheads can be reduced.

Third, every packet requires a link-layer header. For example, every ethernet packet has 18 bytes overhead (2 bytes for type, 6 bytes for source address, 6 bytes for destination address, 4 bytes for CRC). Moreover, an ethernet frame also has 12 bytes of frame spacing and 8 bytes of preamble. Also, packets smaller than 46 bytes have to be padded with additional bytes to make the frame size 64 bytes. Therefore, by combining two 40 bytes packets into a larger packet, 44 bytes are saved. The gains increase further if more than two packets are combined into a larger packet.

For gigabit ethernet, the minimum frame size is 512 bytes, and hence even bigger savings are obtained by combining multiple small packets into a single larger packet. Packet bursting technique reduces the padding overheads by sending back-to-back frames, however, the first frame still incurs the padding

overhead and two consecutive frames must have the required inter-frame spacing.

Fourth, on links such as Ethernets and wireless LANs, terminals need to contend for the channel prior to sending a packet. Reducing the number of 5 packets reduces the contention cost and increases channel utilization.

Finally, aggregation reduces interrupt overheads. Incoming packets generate an interrupt at the receiver. On a modern RISC processor, each interrupt breaks the CPU pipeline, interferes with the cache, requires a change to the memory management context, and forces a substantial number of CPU 10 registers to be saved. An n -fold reduction in packets thus reduces the interrupt and packet overhead by a factor of n .

It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reviewing the above description. The scope of the 15 invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

What is claimed is:

1. A computerized method of reducing packets on a network, the method comprising:
 - 5 combining multiple packets for a destination into an aggregated packet; transmitting the single packet to the destination; separating the aggregated packet into the multiple packets at the destination.
- 10 2. The method of claim 1, wherein a predetermined number of the multiple packets are combined into the aggregated packet.
3. The method of claim 1, wherein the multiple packets received over a predetermined period of time are combined into the aggregated packet.
- 15 4. The method of claim 1 wherein the combining and splitting is implemented in hardware.
5. The method of claim 1 wherein the combining and splitting is implemented 20 in software.
6. The method of claim 1 wherein the destination is an end station.
7. The method of claim 1 wherein the destination is an intermediate station.
- 25 8. The method of claim 1 wherein the combining is performed at a dynamically designated network station.
9. The method of claim 1 wherein the combining is performed at a 30 user-designated network station.

10. A network router comprising:
a routing table to identify a destination for a plurality of packets; and
an aggregator module to combine at least two of the packets with a same destination into an aggregate packet.

5

11. The network router of claim 10, wherein a predetermined number of the multiple packets are combined into the aggregated packet.

10 12. The network router of claim 10, wherein the multiple packets received over a predetermined period of time are combined into the aggregated packet.

13. The network router of claim 10 wherein the aggregator module is a hardware module.

15 14. The network router of claim 10 wherein the aggregator module is a software module.

20 15. A network router comprising :
a routing table to identify a destination for an aggregated packet; and
a splitter module to regenerate a plurality of original packets from a second aggregated packet.

16. The network router of claim 15 wherein the splitter module is a hardware module.

25

17. The network router of claim 15 wherein the splitter module is a software module.

30 18. A computerized system comprising:
an aggregator module to combine multiple packets with a same destination into a first aggregate packet; and

a splitter module to regenerate a plurality of original packets from a second aggregated packet.

19. The computerized system of claim 18, further comprising a signaling
5 mechanism to dynamically establish an aggregator module.

20. A computer readable medium containing executable program instructions
combining network packets, the executable instructions comprising program
instructions for:

10 combining small packets for a same destination into a single larger packet;
adding meta data to the aggregated packet; and
sending the aggregated packet to a receiver.

21. The computer readable medium of claim 20 wherein the combining of
15 small packets into a single larger packet continues until a predetermined number
of small packets have been combined into the larger packet.

22. The computer readable medium of claim 20 wherein the combining of
small packets continues until a predetermined period of time expires.

1 / 4

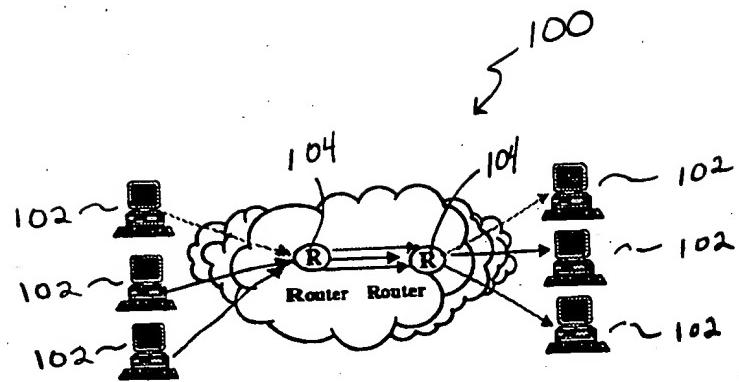


FIG. 1

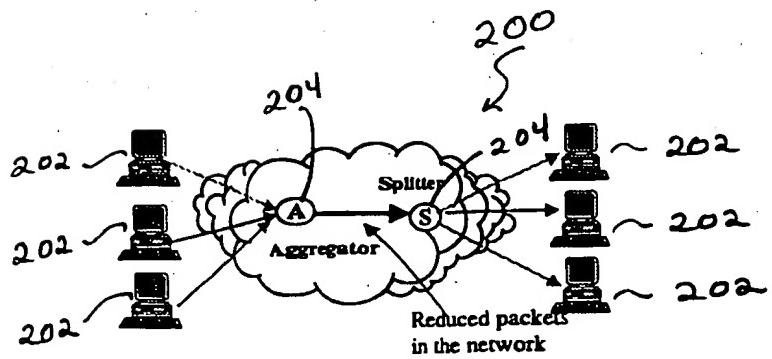


FIG. 2

2 / 4

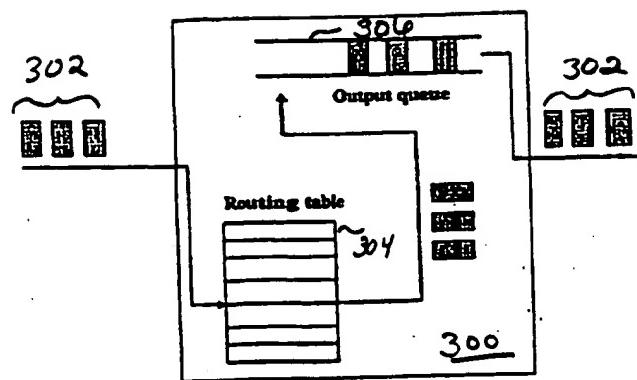


FIG. 3 A

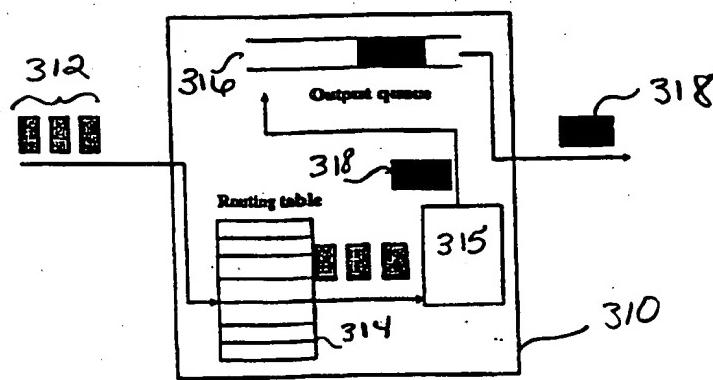


FIG. 3 B

3 / 4

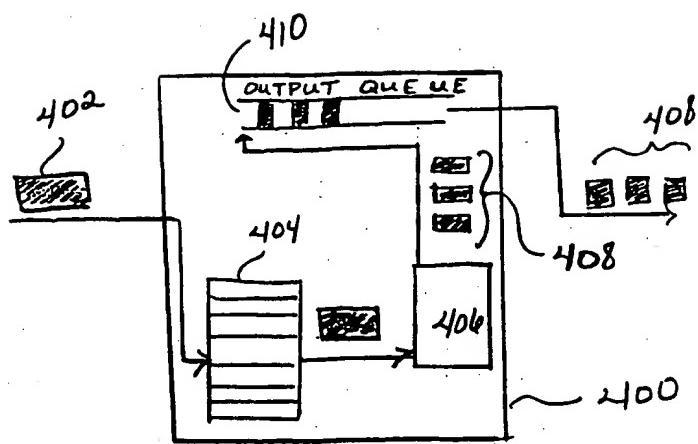
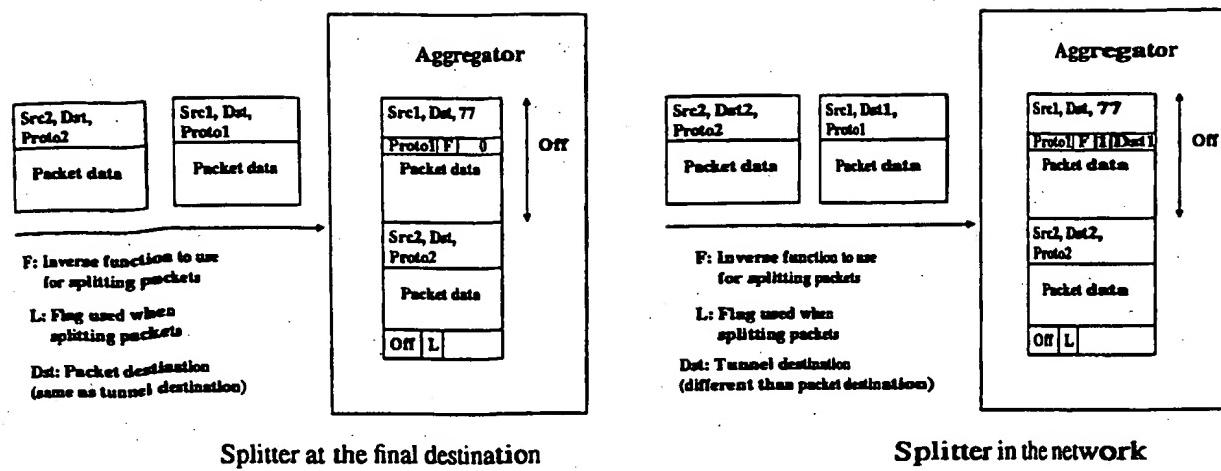


FIG. 4

4 / 4



Splitter at the final destination

Splitter in the network

FIG. 5

INTERNATIONAL SEARCH REPORT

Internal Application No

PCT/US 00/14381

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04L12/56 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	B SUBBIAH ET AL: "User Multiplexing in RTP payload between IP Telephony Gateways" IETF, INTERNET DRAFT, 21 August 1998 (1998-08-21), XP002127741 page 2, line 37 -page 3, line 7 page 4, line 1 -page 5, line 4	1-5, 7, 10-18, 20-22
A	---	6, 8, 9, 19
X	US 5 822 523 A (ROTHSCHILD J J ET AL) 13 October 1998 (1998-10-13) column 8, line 61 -column 9, line 26 column 10, line 20 - line 49 column 14, line 37 - line 51 column 23, line 50 -column 26, line 22 ---	1-3, 5, 6, 10-12, 14, 20-22
	-/-	

Further documents are listed in continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

27 October 2000

06/11/2000

Name and mailing address of the ISA
 European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
 Fax: (+31-70) 340-3016

Authorized officer

Ströbeck, A

INTERNATIONAL SEARCH REPORT

Intern'l Application No

PCT/US 00/14381

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>"BUNDLING SHORT PACKETS IN AN ATM OR PACKET NETWORK"</p> <p>IBM TECHNICAL DISCLOSURE BULLETIN, US, IBM CORP., NEW YORK, vol. 34, no. 48, 1 September 1991 (1991-09-01), pages 360-362, XP000189604 ISSN: 0018-8689 page 360, line 1 - line 14 page 362, line 31 - line 39 ----</p>	1,4-8, 18-20
P,X	<p>WO 00 11849 A (SUBBIAH BARANITHARN) 2 March 2000 (2000-03-02)</p> <p>page 8, line 33 -page 9, line 15 page 10, line 19 -page 11, line 8 page 12, line 19 -page 13, line 15 page 14, line 18 - line 29 -----</p>	1,3,7, 10,12, 15, 18-20,22

INTERNATIONAL SEARCH REPORT
Information on patent family members

International Application No

PCT/US 00/14381

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
US 5822523	A 13-10-1998	AU 1747197 A		22-08-1997
		EP 1012724 A		28-06-2000
		JP 2000504133 T		04-04-2000
		WO 9728502 A		07-08-1997
		US 6018766 A		25-01-2000
WO 0011849	A 02-03-2000	AU 5771199 A		14-03-2000

CORRECTED VERSION(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
30 November 2000 (30.11.2000)

PCT

(10) International Publication Number
WO 00/72532 A1(51) International Patent Classification⁷: H04L 12/56,
29/06Court, Somerset, NJ 08873 (US). SUDAME, Predeep
[US/US]; 18J Reler Lane, Somerset, NJ 08873 (US).

(21) International Application Number: PCT/US00/14381

(74) Agent: VIKSNINS, Ann, S.; Schwegman, Lundberg,
Woessner & Kluth, P.O. Box 2938, Minneapolis, MN
55402 (US).

(22) International Filing Date: 24 May 2000 (24.05.2000)

(25) Filing Language: English

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE,
DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU,
ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS,
LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ,
PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT,
TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(26) Publication Language: English

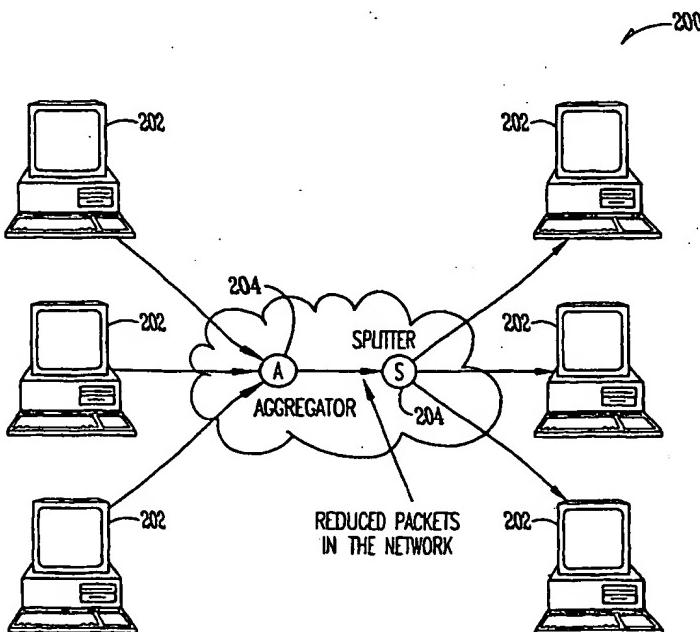
(30) Priority Data:
60/135,605 24 May 1999 (24.05.1999) US(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European
patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,
IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG,
CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).(71) Applicant: RUTGERS, THE STATE UNIVERSITY
OF NEW JERSEY [US/US]; Old Queens Building,
Somerset and George Streets, New Brunswick, NJ 08903
(US).

(71) Applicants and

(72) Inventors: BADRINATH, B.R. [US/US]; 4 Depalma

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR NETWORK PACKET REDUCTION



WO 00/72532 A1

(57) **Abstract:** A system of and method for reducing network packet traffic is described. Embodiments of the present invention combine small packets with the same destination into a larger packet, thereby reducing the number of packets flowing in the network. Small packets dominate Internet traffic patterns today. The overhead due to queuing, processing, routing, differential treatment, and delivery is the same irrespective of the size of the packet. Hence, small packets impose a relatively higher overhead on the network. The domination of small packet traffic is expected to get only worse with the ever increasing web applications and the introduction of new applications such as Voiceover IP.

**Published:**

— *with international search report*

(48) Date of publication of this corrected version:

18 April 2002

(15) Information about Correction:

see PCT Gazette No. 16/2002 of 18 April 2002, Section II

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

System and Method for Network Packet Reduction

5

Field of the Invention

The present invention relates generally to computer networks and more specifically to reducing packets on a network.

Copyright Notice/Permission

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever. The following notice applies to the software and data as described below and in the drawings hereto: Copyright © 2000 Rutgers, 10 The State University of New Jersey, All Rights Reserved.
15

Background of the Invention

Data communication in a computer network involves the exchange of data between two or more entities interconnected by communication links and subnetworks. These entities are typically software programs executing on hardware computer platforms, such as end stations and intermediate stations. Examples of an intermediate station may be a router or switch which interconnects the communication links and subnetworks to enable transmission of data between the end stations. A local area network (LAN) is an example of a subnetwork that provides relatively short distance communication among the interconnected stations; in contrast, a wide area network (WAN) enables long distance communication over links provided by public or private telecommunications facilities. Communication software executing on the end stations correlates and manages data communication with other end stations. The stations typically communicate by exchanging discrete 20
25
30

packets or frames of data according to predefined protocols.

Collectively, these hardware and software components comprise a communications network.

Small packets (less than the default link MTU of 512 bytes) dominate Internet traffic patterns today. Wide-area Internet traffic measurements show that as many as sixty percent (60%) of the packets transmitted over the Internet backbone are less than 44 bytes. Web servers, for example, are a leading generator of small packets. All of the TCP ACKs generated by the clients have the same destination (the web server). These ACKs are just 40 bytes in size. Another example of small packet traffic is ICP (Internet Cache Protocol) queries generated among web proxies. ICP queries are about 60 bytes in size. This domination of small packets in Internet traffic patterns is expected to get only worse with the ever increasing web traffic and with introduction of new applications such as VoIP (Voiceover-IP applications generate small audio samples, about 40 to 60 byte packets every 40 ms). The network overhead cost due to queuing, processing, routing, differential treatment and delivery is the same irrespective of the size of the packet. Hence, small packets impose a relatively higher overhead on the network.

Therefore, in order to improve overall network performance, there is a need to reduce the number of packets in the network.

Summary of the Invention

A system of and method for reducing network packet traffic is described. Embodiments of the present invention combine small packets with the same destination into a larger packet, thereby reducing the number of packets flowing in the network. Small packets dominate Internet traffic patterns today. The overhead due to queuing, processing, routing, differential treatment, and delivery is the same irrespective of the size of the packet. Hence, small packets impose a relatively higher overhead on the network. The domination of small packet traffic is expected to get only worse with the ever increasing web applications and

the introduction of new applications such as Voiceover IP. To mitigate the impact of small packets, a novel approach of combining multiple small packets going towards the same destination is provided.

One aspect of the invention is a computerized method of reducing 5 packets on a network. The method comprises combining multiple packets for a destination into an aggregated packet, and separating the aggregated packet into the multiple packets at the destination.

Still other and further embodiments, aspects and advantages of the 10 invention will become apparent by reference to the drawings and by reading the following detailed description.

Brief Description of the Drawings

Figure 1 is a block diagram of a traditional communication network.

Figure 2 is a block diagram of a communication network according to an example embodiment of the invention.

15 Figure 3A is a block diagram of a traditional router.

Figure 3B is a block diagram of an aggregator, such as the aggregator shown in Figure 2, according to an example embodiment of the present invention.

20 Figure 4 is a block diagram of a splitter, such as the splitter shown in Figure 2, according to an example embodiment of the present invention.

Figure 5 is a block diagram showing how metadata is stored when two small packets are combined at a router according to an example embodiment of the invention.

Description of the Embodiments

25 In the following detailed description of the embodiments, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention.
30 It is to be understood that other embodiments may be utilized and that structural, logical and electrical changes may be made without departing

from the scope of the present inventions. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present inventions is defined only by the appended claims.

A system of and method for reducing network packet traffic is described below. Embodiments of the present invention combine small packets with the same destination into a larger packet, thereby reducing the number of packets flowing in the network.

Figure 1 is a block diagram of a traditional communication network 100. As shown in Figure 1, the traditional communication network 100 comprises a plurality of end stations 102 and a plurality of intermediate stations 104 interconnected by communications links. In the example shown in Figure 1, the intermediate links are routers 104.

Figure 2 is a block diagram of a communication network 200 according to an example embodiment of the present invention. As shown in Figure 2, the communication network 200 comprises a plurality of end stations 202 and a plurality of intermediate stations 204. The intermediate stations 204 include an aggregator and a splitter which are further described below. The aggregator and the splitter function to reduce the packets on the communications link between them. The aggregator combines a small packets into a larger packet so that the routers see a smaller number of packets. The splitter separates these packets so that the subsequent stations in the network are unaware of the packet aggregation. Aggregators, as well as splitters, can be placed at intermediate stations or at end stations in the network. The placement of aggregators and splitters may be varied depending on traffic patterns, level of congestion and software application needs.

In one embodiment, the aggregators as well as the splitters can be the end stations in the network. In an alternate embodiment, the aggregators and splitters are placed at ingress and egress routers in the network. The placement of aggregators and splitters may be manually configured or may be dynamically assigned according to a signaling mechanism of the

present invention which is further described below. In another embodiment, the aggregated packets on the communications links between the aggregator and the splitter are treated as a special class of traffic to further enhance network performance. The number of packets in the network are reduced by combining small packets multiple times at various points in the network.

In one embodiment, the placement of the aggregator in the network is based on the application and is dynamically determined by the receiver of the packets. A signaling mechanism of the present invention allows the receiver of the packets to place aggregators at points in the network where packets can be combined. If the aggregation at that location in the network is ineffective, the aggregators can be dynamically or manually removed and, if desired, repositioned at another point in the network. In other words, if the application ceases to need aggregation, the network can remove the aggregators without affecting the normal network routing. Thus, the signalling mechanism allows the aggregation of packets to be initiated and stopped dynamically. In an alternate embodiment, the signaling mechanism can be initiated and stopped by manual configure according to a user.

Figure 3A is a block diagram of a traditional router 300. As shown in Figure 3A, three packets 302 are received by router 300. The router 300 performs a routing table 304 lookup to determine the output queue 306 for each one of the packets. All three packets 302 leave router 300 and are transmitted across a communication link.

Figure 3B is a block diagram of an aggregator such as the aggregator shown in Figure 2. As used herein, an aggregator refers to any end station or intermediate station capable of combining multiple small packets to the same destination and sending them as one packet. For example, an aggregator may be implemented as a hardware module or a software module in a router. As shown in Figure 3B, three packets 312 are received by router 310. The router 310 performs a routing table 304 lookup to determine the destination of the packets

302. However, unlike the traditional router shown in Figure 3A, the router 310 comprises an aggregation module 315. In this example, the three packets 302 have the same destination. The aggregation module 315 combines the three packets 312 into a single packet 318 referred to herein as an "aggregated packet." 5 The aggregated packet 318 is placed in the output queue 316 for the destination and unlike the traditional router, only the single aggregated packet leaves the router 310 and is transmitted across the communication link.

Once such an aggregated packet reaches the destination, all the original packets can be regenerated. The node that separates aggregated packets is 10 referred to herein as a "splitter." Figure 4 is a block diagram of a splitter 400. As used herein, a splitter refers to any end station or intermediate station capable of separating an aggregated packets into the original multiple packets. For example, a splitter may be implemented as a hardware module or a software module in a router. As shown in Figure 4, an aggregated packet 402 is received 15 by router 400. The router 400 performs a routing table 404 lookup to determine the destination of the packet 402. However, unlike the traditional router shown in Figure 3A, the router 400 comprises a splitter module 406. The splitter module 400 separates the aggregated packet 406 into the original multiple packets 408. In one embodiment, the multiple packets 408 are placed in the 20 output queue 410 for router 400.

Aggregation as well as splitting can be performed either at layer 2 (MPLS) or at layer 3 (network layer) of the protocol stack. Aggregators at layer 2 require a hardware implementation. In an example embodiment, aggregators and splitters are implemented using transformer tunnels. Transformer tunnels are 25 described in more detail in Predeep Sudame and B. R. Bardinath, "Transformer tunnels: A framework for providing route-specific adaptions," in *Proceedings of the USENIX Annual Technical Conference*, June 1998, pp. 191-200, which is herein incorporated by reference. However, embodiments of the present invention are not limited to being implemented with transformer tunnels and 30 alternate embodiments are contemplated. To control the number of routing table entries at an aggregator node, the splitter can be a node other than the final

destination. This way, packets belonging to multiple destinations (for example, a subnet) can be combined using a single tunnel. Aggregators may also be dynamically placed in the network so that the bottleneck links lie between the splitters and the aggregators.

5 A computerized method of reducing packets on a network according an example embodiment of the invention comprises combining multiple packets for a destination into an aggregated packet, transmitting the single packet to the destination, and separating the aggregated packet into the multiple packets at the destination. In one embodiment, an aggregation function combines the multiple 10 packets for a destination into an aggregated packet using a mechanism similar to TCP delayed ACKs. Every small packet is delayed by a small amount of time. The process of combining small packets continues until either m (a parameter to the aggregation function described below) packets have been combined, or a timeout occurs. If no other small packets arrive before the timeout, the delayed 15 packet is sent as is. This aggregation mechanism is different from IP reassembly (performed after IP fragmentation). An aggregation function combines small packets, along with their IP headers, to get a larger packet. The receiver regenerates all the original packets when it receives such an aggregated packet. Thus, this mechanism can be used even for protocols that honor message 20 boundaries (e.g., UDP).

In one embodiment, for the receiver to be able to regenerate the original packets, some information is added (metadata) to the aggregated packet. Figure 5 is a block diagram showing how metadata is stored when two small packets are combined at a router. The generated larger packet has a new protocol number 25 (77) in the IP header. (The original protocol therefore has to be stored as part of the metadata, along with the original destination of the packet, if it is different than the address of the splitter node.) The receiver of this packet handles the packet to the protocol handler for this new protocol. This handler is responsible for removing all the metadata and restoring the packets.

30 As small packets are delayed for some time in the hope that they can be combined with other small packets, the delayed packets see increased latency.

The latency depends on the packet inter-arrival rate, and the timeout set by the applications. In the best case, when m packets are combined, the timeout does not come into picture, because the buffers are flushed immediately when the packets are combined.

5 An example embodiment of an aggregation function as described above is provided below. The aggregation function is provided for illustrative purposes only.

Parameters:

10 timeout, m : max packets to be combined

Initialization code:

 delayed_packet = null;

15 aggregate (p) {
 if (p is large) {
 cancel timer; send delayed_packet;
 delayed_packet=null;
 }
 20 send p; return PACKET_SENT;
 }
 if (delayed_packet exists) {
 if (p and delayed_packet can be combined) {
 delayed_packet = combine (delayed_packet, p);
 if (m packets have been combined) {
 cancel timer; send delayed_packet;
 delayed_packet = null;
 }
 return PACKET_SENT;
 }
 30 } else {
 cancel timer; send delayed_packet;
 set timer; delayed_packet = p;
 return PACKET_SENT;
 }
 35 } else {
 set timer; delayed_packet = p;
 return PACKET_SENT;
 }
 }
40 Timer expires:
 send delayed_packet; delayed_packet = null;

- Given that there is sufficient evidence of small packet traffic in the Internet, there are a number of benefits to be obtained by packet-level aggregation. First, most router implementations (such as FreeBSD, Linux, and many Cisco routers) are packet limited rather than byte limited. They place a limit on the number of packets that the router can queue irrespective of the size of each packet. For reasons of efficiency, routers store packet headers (whose sizes are independent of the packet sizes) in fast memory, and the packet payload in slow memory. The fast memory being expensive decides the queue size. Every packet, irrespective of the size, consumes one slot in the queue.
- Combining small packets into a larger packet reduces the number of packets seen at routers, and hence helps absorb bursts (reduce losses) for a given level of congestion.
- Second, every packet at a router requires a routing table lookup and incurs the same performance cost irrespective of the size of the packet. Combining small packets at a router reduces the number of packets seen by all the subsequent routers. Moreover, to implement a specific per-hop behavior (as in differential services), routers have to give preferential treatment to every packet of a particular class. By combining small packets belonging to the same class, these overheads can be reduced.
- Third, every packet requires a link-layer header. For example, every ethernet packet has 18 bytes overhead (2 bytes for type, 6 bytes for source address, 6 bytes for destination address, 4 bytes for CRC). Moreover, an ethernet frame also has 12 bytes of frame spacing and 8 bytes of preamble. Also, packets smaller than 46 bytes have to be padded with additional bytes to make the frame size 64 bytes. Therefore, by combining two 40 bytes packets into a larger packet, 44 bytes are saved. The gains increase further if more than two packets are combined into a larger packet.

For gigabit ethernet, the minimum frame size is 512 bytes, and hence even bigger savings are obtained by combining multiple small packets into a single larger packet. Packet bursting technique reduces the padding overheads by sending back-to-back frames, however, the first frame still incurs the padding

overhead and two consecutive frames must have the required inter-frame spacing.

Fourth, on links such as Ethernets and wireless LANs, terminals need to contend for the channel prior to sending a packet. Reducing the number of 5 packets reduces the contention cost and increases channel utilization.

Finally, aggregation reduces interrupt overheads. Incoming packets generate an interrupt at the receiver. On a modern RISC processor, each interrupt breaks the CPU pipeline, interferes with the cache, requires a change to the memory management context, and forces a substantial number of CPU 10 registers to be saved. An n -fold reduction in packets thus reduces the interrupt and packet overhead by a factor of n .

It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reviewing the above description. The scope of the 15 invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

What is claimed is:

1. A computerized method of reducing packets on a network, the method comprising:
 - 5 combining multiple packets for a destination into an aggregated packet; transmitting the single packet to the destination; separating the aggregated packet into the multiple packets at the destination.
- 10 2. The method of claim 1, wherein a predetermined number of the multiple packets are combined into the aggregated packet.
- 15 3. The method of claim 1, wherein the multiple packets received over a predetermined period of time are combined into the aggregated packet.
4. The method of claim 1 wherein the combining and splitting is implemented in hardware.
5. The method of claim 1 wherein the combining and splitting is implemented 20 in software.
6. The method of claim 1 wherein the destination is an end station.
7. The method of claim 1 wherein the destination is an intermediate station.
- 25 8. The method of claim 1 wherein the combining is performed at a dynamically designated network station.
9. The method of claim 1 wherein the combining is performed at a 30 user-designated network station.

10. A network router comprising:
a routing table to identify a destination for a plurality of packets; and
an aggregator module to combine at least two of the packets with a same destination into an aggregate packet.

5

11. The network router of claim 10, wherein a predetermined number of the multiple packets are combined into the aggregated packet.

- 10 12. The network router of claim 10, wherein the multiple packets received over a predetermined period of time are combined into the aggregated packet.

13. The network router of claim 10 wherein the aggregator module is a hardware module.

- 15 14. The network router of claim 10 wherein the aggregator module is a software module.

15. A network router comprising :
a routing table to identify a destination for an aggregated packet; and
20 a splitter module to regenerate a plurality of original packets from a second aggregated packet.

16. The network router of claim 15 wherein the splitter module is a hardware module.

25

17. The network router of claim 15 wherein the splitter module is a software module.

18. A computerized system comprising:
30 an aggregator module to combine multiple packets with a same destination into a first aggregate packet; and

a splitter module to regenerate a plurality of original packets from a second aggregated packet.

19. The computerized system of claim 18, further comprising a signaling mechanism to dynamically establish an aggregator module.

20. A computer readable medium containing executable program instructions combining network packets, the executable instructions comprising program instructions for:

10 combining small packets for a same destination into a single larger packet; adding meta data to the aggregated packet; and sending the aggregated packet to a receiver.

21. The computer readable medium of claim 20 wherein the combining of 15 small packets into a single larger packet continues until a predetermined number of small packets have been combined into the larger packet.

22. The computer readable medium of claim 20 wherein the combining of small packets continues until a predetermined period of time expires.

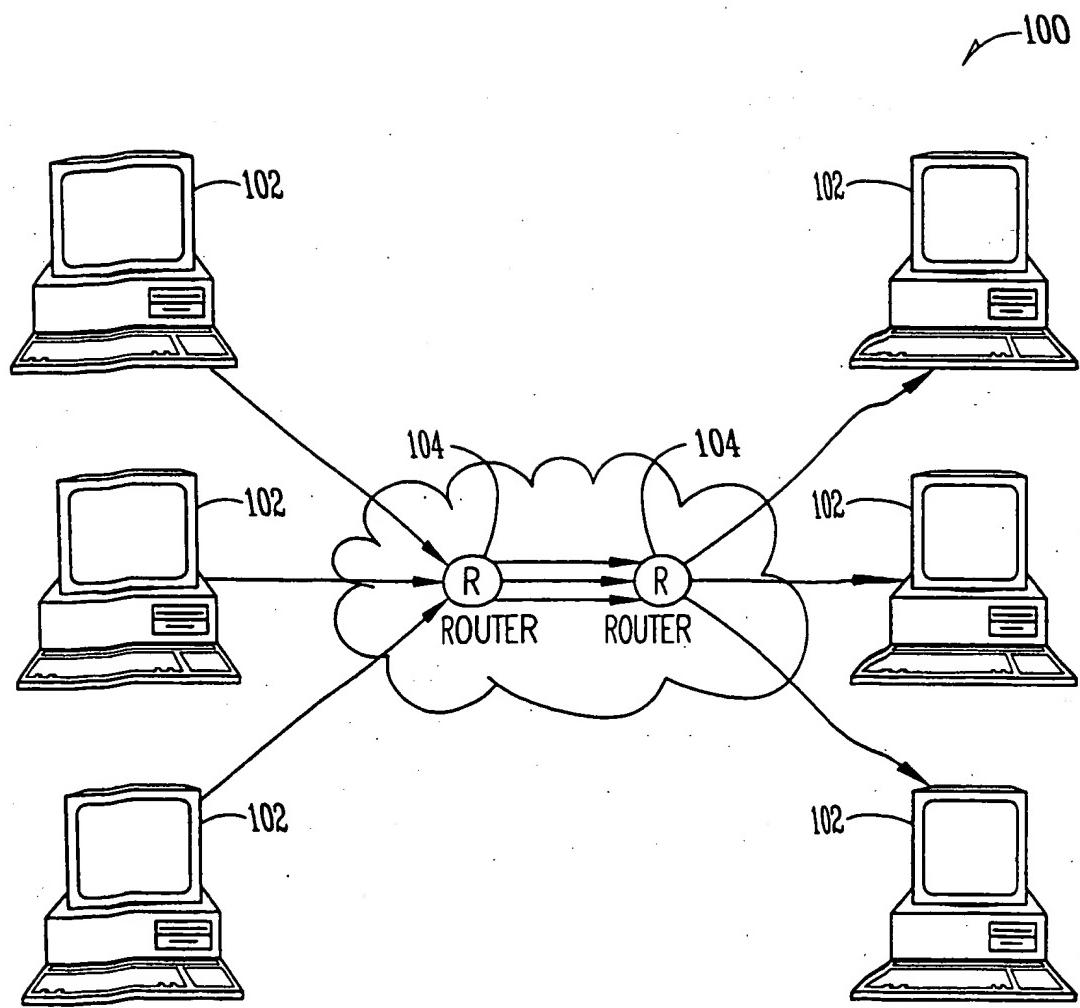


Fig. 1

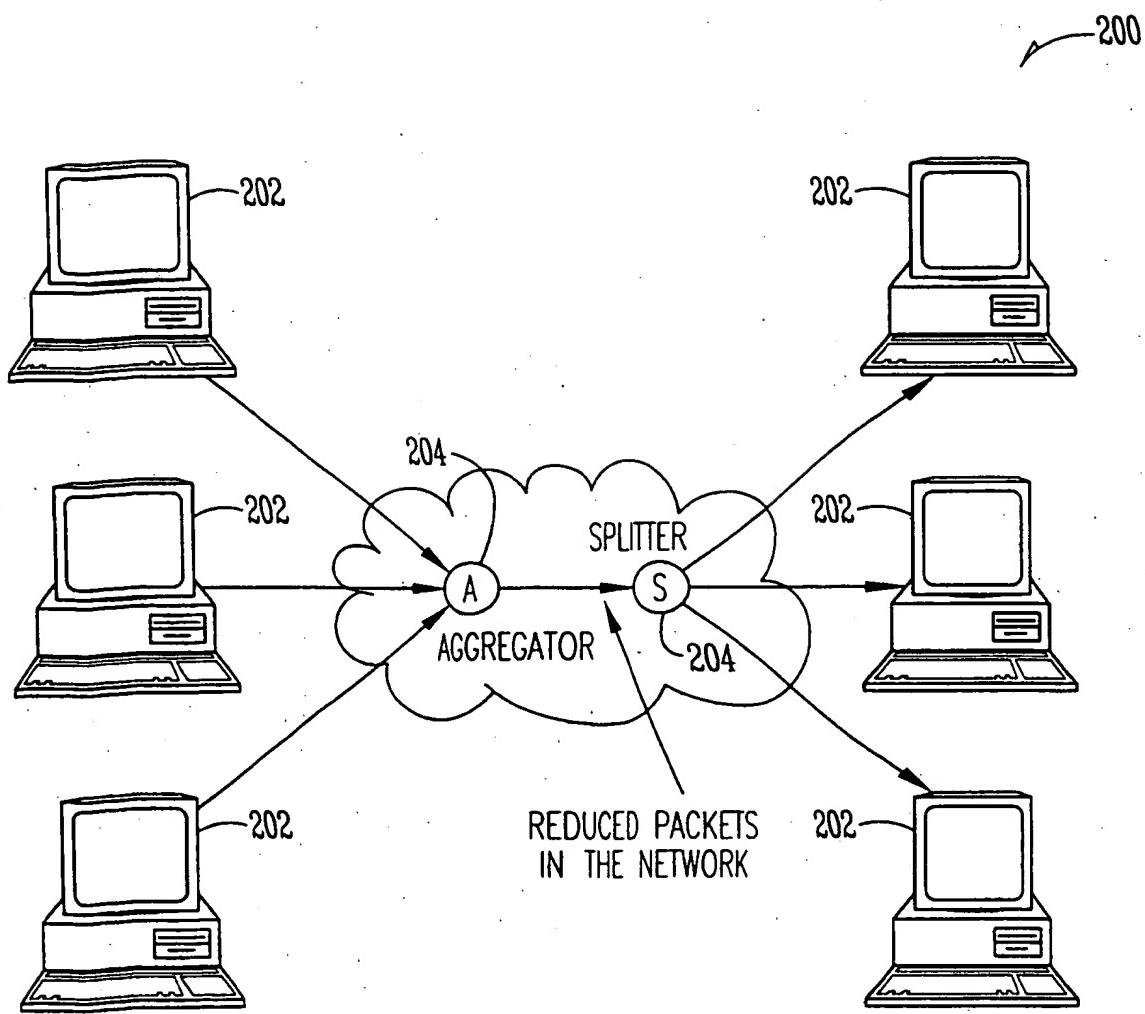


Fig.2

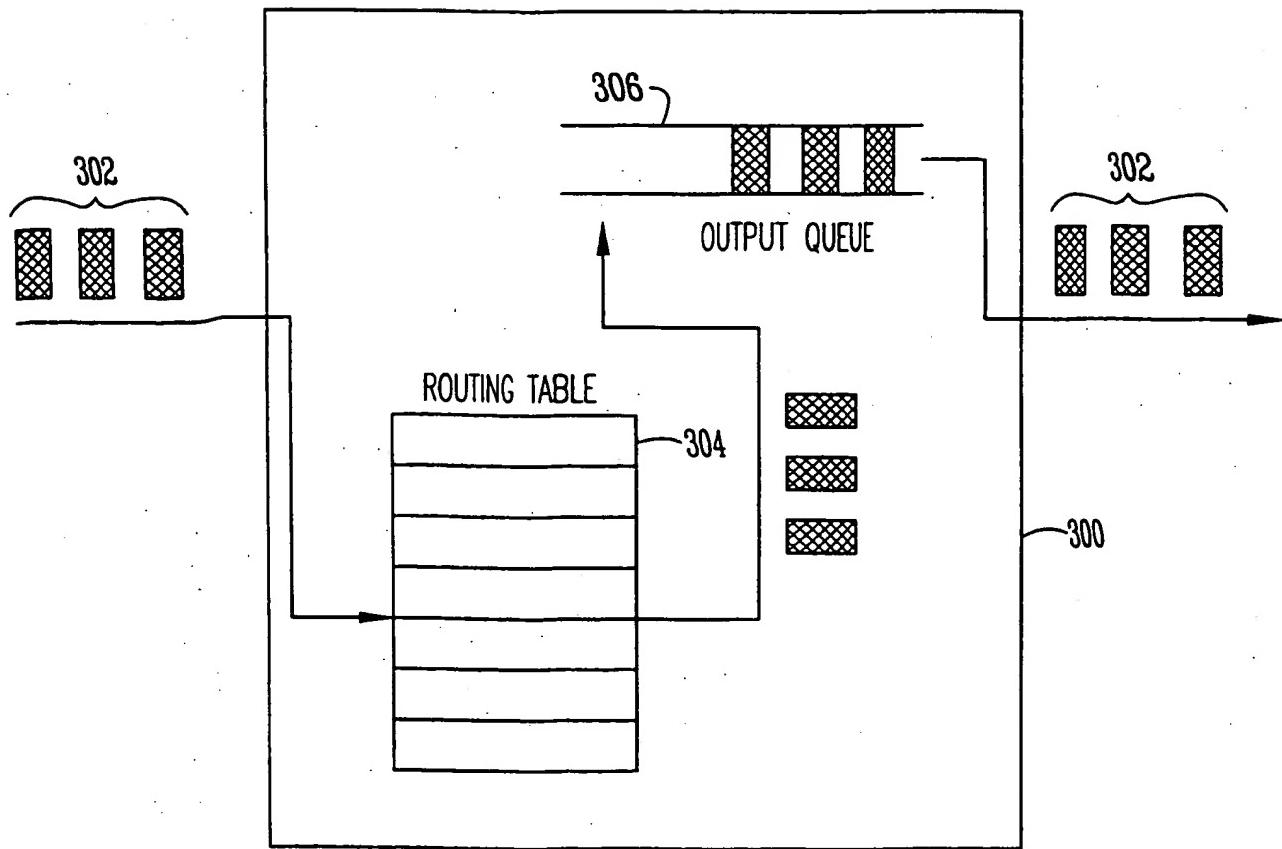


Fig. 3A

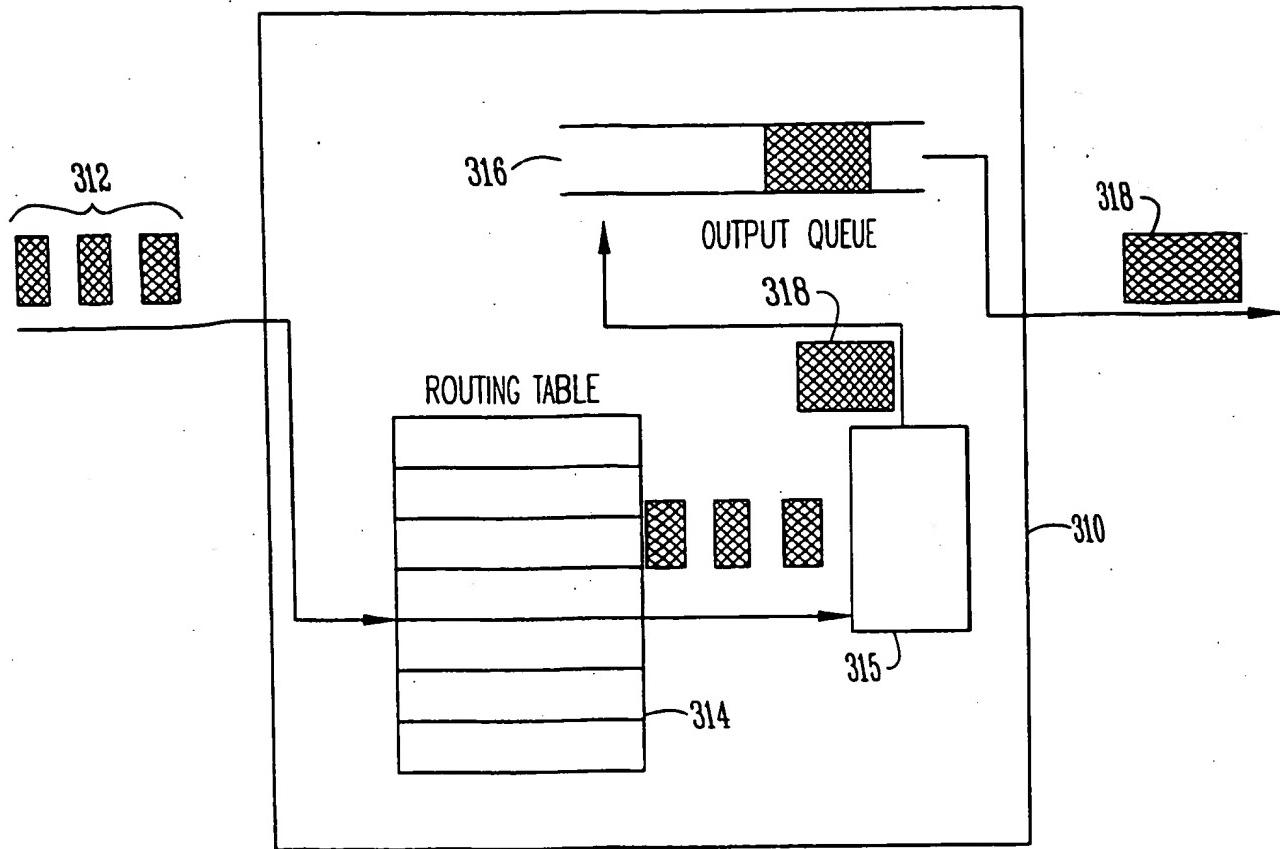


Fig. 3B

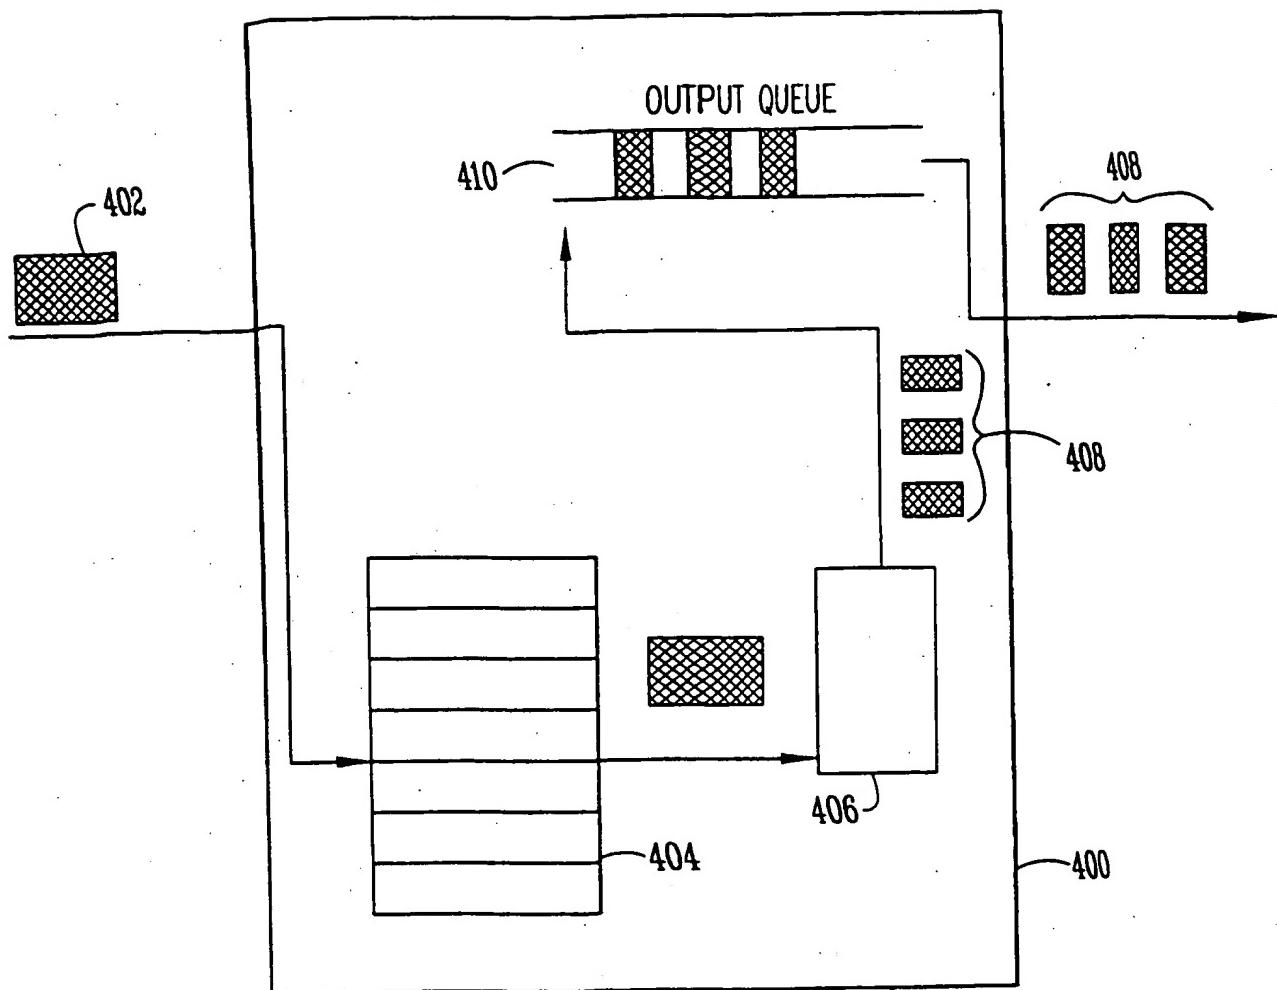
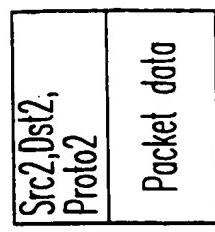
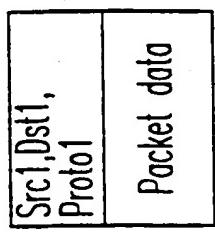
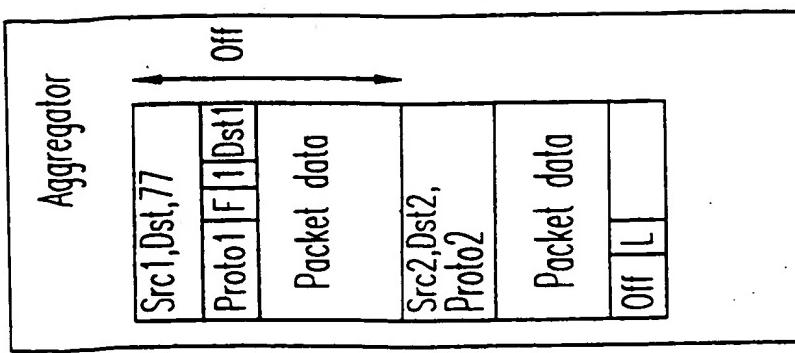


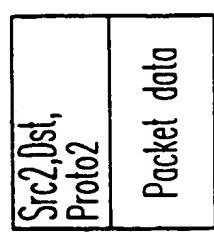
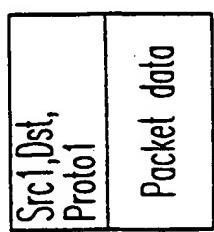
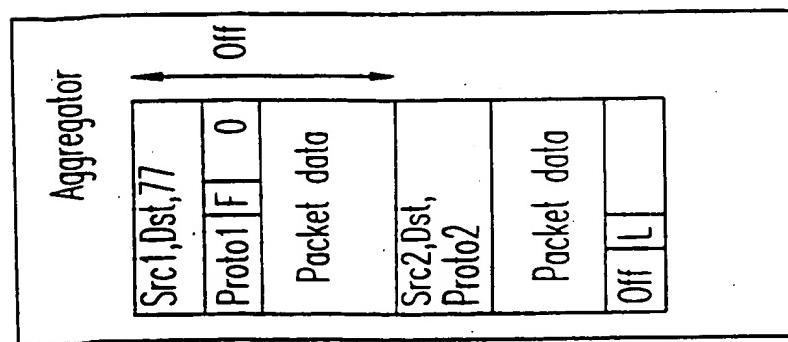
Fig. 4



F: Inverse function to use
for splitting packets

L: Flag used when
splitting packets

Dst: Tunnel destination
(different than packet destination)



F: Inverse function to use
for splitting packets

L: Flag used when
splitting packets

Dst: Packet destination
(same as tunnel destination)

Splitter at the final destination

Splitter in the network

Fig. 5

INTERNATIONAL SEARCH REPORT

Internal Application No

PCT/US 00/14381

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04L12/56 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	B SUBBIAH ET AL: "User Multiplexing in RTP payload between IP Telephony Gateways" IETF, INTERNET DRAFT, 21 August 1998 (1998-08-21), XP002127741 page 2, line 37 -page 3, line 7 page 4, line 1 -page 5, line 4	1-5,7, 10-18, 20-22
A	---	6,8,9,19
X	US 5 822 523 A (ROTHSCHILD J J ET AL) 13 October 1998 (1998-10-13) column 8, line 61 -column 9, line 26 column 10, line 20 - line 49 column 14, line 37 - line 51 column 23, line 50 -column 26, line 22 ---	1-3,5,6, 10-12, 14,20-22

 Further documents are listed in the continuation of box C. Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

27 October 2000

06/11/2000

Name and mailing address of the ISA
 European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
 Fax: (+31-70) 340-3016

Authorized officer

Ströbeck, A

INTERNATIONAL SEARCH REPORT

Intern'l Application No

PCT/US 00/14381

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>"BUNDLING SHORT PACKETS IN AN ATM OR PACKET NETWORK"</p> <p>IBM TECHNICAL DISCLOSURE BULLETIN, US, IBM CORP. NEW YORK, vol. 34, no. 4B, 1 September 1991 (1991-09-01), pages 360-362, XP000189604 ISSN: 0018-8689 page 360, line 1 - line 14 page 362, line 31 - line 39 -----</p>	1,4-8, 18-20
P,X	<p>WO 00 11849 A (SUBBIAH BARANITHARN) 2 March 2000 (2000-03-02)</p> <p>page 8, line 33 -page 9, line 15 page 10, line 19 -page 11, line 8 page 12, line 19 -page 13, line 15 page 14, line 18 - line 29 -----</p>	1,3,7, 10,12, 15, 18-20,22

INTERNATIONAL SEARCH REPORT

Information on patent family members

Internat'l Application No

PCT/US 00/14381

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
US 5822523	A 13-10-1998	AU 1747197 A		22-08-1997
		EP 1012724 A		28-06-2000
		JP 2000504133 T		04-04-2000
		WO 9728502 A		07-08-1997
		US 6018766 A		25-01-2000
-----	-----	-----	-----	-----
WO 0011849	A 02-03-2000	AU 5771199 A		14-03-2000
-----	-----	-----	-----	-----

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.